

The Impact of Human–Automation Collaboration in Decentralized Multiple Unmanned Vehicle Control

By M. L. CUMMINGS, *Senior Member IEEE*, JONATHAN P. HOW, *Senior Member IEEE*,
ANDREW WHITTEN, AND OLIVIER TOUPET

ABSTRACT | For future systems that require one or a small team of operators to supervise a network of automated agents, automated planners are critical since they are faster than humans for path planning and resource allocation in multivariate, dynamic, time-pressured environments. However, such planners can be brittle and unable to respond to emergent events. Human operators can aid such systems by bringing their knowledge-based reasoning and experience to bear. Given a decentralized task planner and a goal-based operator interface for a network of unmanned vehicles in a search, track, and neutralize mission, we demonstrate with a human-on-the-loop experiment that humans guiding these decentralized planners improved system performance by up to 50%. However, those tasks that required precise and rapid calculations were not significantly improved with human aid. Thus, there is a shared

space in such complex missions for human-automation collaboration.

KEYWORDS | Command and control; decentralized task planning; decision support systems; human-automation interaction; human supervisory control; unmanned vehicles

I. INTRODUCTION

The use of unmanned vehicles (UVs) has recently revolutionized U.S. military operations. In the past year, the U.S. Air Force reached an important milestone in that it now has more unmanned aerial vehicles (UAVs) than manned aircraft. Other countries are following this lead, with recent significant advances in UAV heavy-lift capacity by Israel and the United Kingdom. These dramatic advances are not only limited to the military as the international civil sector also looks to such unmanned technologies to aid operations such as fighting forest fires, undersea exploration, monitoring wildlife, inspecting bridges, and supporting first responders such as police and rescue organizations. UAV expenditures alone are predicted to more than double in the next ten years, and are expected to exceed \$80 billion [1].

Accompanying these impressive advances in UV platforms are equally dramatic leaps in intelligent UV control

Manuscript received July 16, 2010; revised June 7, 2011; accepted October 7, 2011. This work was supported by a U.S. Office of Naval Research STTR under the guidance of Aurora Flight Sciences.

M. L. Cummings and **J. P. How** are with the Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: MissyC@mit.edu; jhow@mit.edu).

A. Whitten is with the United States Air Force, Edwards Air Force Base (AFB), CA 93524 USA (e-mail: andrew.whitten@edwards.af.mil).

O. Toupet was with the Research and Development Center, Aurora Flight Sciences, Cambridge, MA 02142 USA (e-mail: otoupet@gmail.com).

Digital Object Identifier: 10.1109/JPROC.2011.2174104

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2011	2. REPORT TYPE		3. DATES COVERED 00-00-2011 to 00-00-2011		
4. TITLE AND SUBTITLE The Impact of Human-Automation Collaboration in Decentralized Multiple Unmanned Vehicle Control			5a. CONTRACT NUMBER		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Massachusetts Institute of Technology,Cambridge,MA,02139			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT For future systems that require one or a small team of operators to supervise a network of automated agents automated planners are critical since they are faster than humans for path planning and resource allocation in multivariate dynamic, time-pressured environments. However, such planners can be brittle and unable to respond to emergent events. Human operators can aid such systems by bringing their knowledge-based reasoning and experience to bear. Given a decentralized task planner and a goal-based operator interface for a network of unmanned vehicles in a search, track, and neutralize mission, we demonstrate with a human-on-the-loop experiment that humans guiding these decentralized planners improved system performance by up to 50%. However, those tasks that required precise and rapid calculations were not significantly improved with human aid. Thus, there is a shared space in such complex missions for human?automation collaboration.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 12	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

systems. Point-and-click UAV operations are routine such that operators only have to click a place on a map to command a vehicle, and the onboard autonomy determines the most appropriate control actions. Such advances over legacy systems that require typical stick-and-rudder skills have revolutionized military operations in that traditional pilots are no longer needed to control such systems. Indeed, while not yet operational, research and development across a number of laboratories has shown that networks of UAVs, unmanned ground vehicles, and even other UV types such as unmanned underwater vehicles and unmanned surface vessels can be controlled by a single operator [2]–[5]. This single-operator, multiple-unmanned-vehicle architecture is the crux of the Department of Defense’s vision for network-centric operations, where a network of intelligent UVs is able to semiautonomously work with a small group of human operators to execute dynamic missions in time-critical scenarios [6].

While all of the previously mentioned studies examined the ability of underlying automation (in the form of planning and control algorithms) to control a network of heterogeneous unmanned vehicles (UxVs), a significant limitation of this work is a lack of investigation of critical human–automation collaboration issues. In all future visions of operating networks of UVs, humans are expected to assume some kind of supervisory role, through high-level goal expressions which could include resource allocation, target designation, approval for weapons release, etc. While research has shown that one operator can theoretically control multiple UVs with varying levels of embedded autonomy [7]–[10], there have been no previous studies that examine how an operator interacts with decentralized UV planners. Moreover, whether such human interaction can improve, or possibly degrade, overall system performance has not previously been addressed. This paper fills this gap by detailing an experiment conducted to investigate what impact a human operator has on the overall performance of a decentralized UV network, and how such interactions could be improved to enhance overall human–system performance.

II. DECENTRALIZED PLANNING FOR MULTIPLE UV MANAGEMENT

We developed the Onboard Planning System for UxVs Supporting Expeditionary Reconnaissance and Surveillance (OPS-USERS) system to provide a planning framework for a team of autonomous agents under human supervisory control participating in expeditionary missions, which rely heavily on intelligence, surveillance, and reconnaissance. The mission environment contains an unknown number of mobile targets, each of which may be friendly, hostile, or unknown. The mission scenario is multiobjective, and includes finding as many targets as possible, keeping accurate position estimates of hostile and unknown targets, and neutralizing all hostile targets. It is

assumed that static features in the environment, such as terrain type, are known *a priori*, but dynamic features, such as target locations, are not.

A. Architecture

The OPS-USERS system architecture is specifically designed to meet the challenges associated with an automated decision-making system integrated with a human operator on the loop. Two key challenges are: 1) balancing the roles and responsibilities of the human operator and the automated planner, and 2) optimizing resource allocation to accomplish each of the mission objectives. The system relies on the relative strengths of both humans and automation in that a human operator provides valuable judgment and field experience in recognizing patterns and defining new goals as environmental changes dictate. A significant benefit of automation is its ability to provide raw computational power and rapid optimization capability on the order of seconds for resource allocation problems that take human decision makers minutes or hours to solve.

In OPS-USERS, decision making responsibility is layered to promote goal-based reasoning such that the human guides the autonomy, but the automation assumes the bulk of the computation. The automated planner is responsible for decisions requiring rapid calculations or optimization, and the human operator supervises the planner for high-level goals such as where to focus the search and which tasks should be included in the overall plan, as well as those tasks that require strict human approval, such as weapons release. Fig. 1 depicts the role allocation balance between the human and automation.

In this goal-based system, the operator is responsible for strategic decision making that includes prioritizing which tasks should be performed at the current stage in the mission. The operator interface that allows such interactions is discussed in detail in Section II-F. The automated planner is responsible for tactical planning by deciding which UxVs perform which task, and forming a task schedule for each UxV. Together, the human operator and the automation work jointly to define plans that meet operational goals. For example, the automation helps the

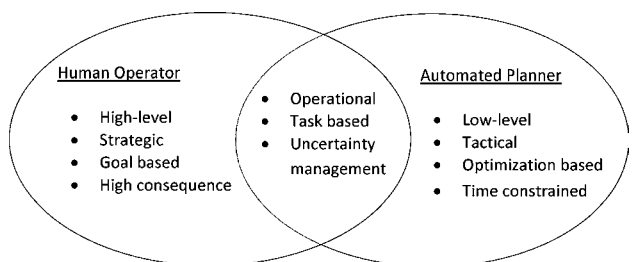


Fig. 1. Decision making allocation between human operator and automated planner.

human by estimating when and where the UxVs should revisit targets based on observed target behavior and sensing capabilities onboard the UxVs, such as present day motion target indicators. The human helps the automation by guiding it to areas the operator believes are more likely to contain undetected targets. Thus, this overall framework allows the human to make global decisions regarding locations and targets of interest, while the automated planner optimizes the local trajectory of each UxV to maximize searching and tracking.

The autonomy only manages uncertainty related to the targets, which is modeled by spatial probability distributions and visually represented to operators as grayed out regions on the interface. The operator aids in target uncertainty management by reasoning about where unfound or lost targets may hide and create search tasks to guide the UxVs to these areas. The autonomy maintains these distributions over time, updating them as the UxVs gather new information and locally optimizes the UxV paths (within a limited planning horizon due to computational constraints). The uncertainty related to the predicted future location of the targets is also managed by the autonomy through periodic revisits. There is, however, no uncertainty in the planning itself, i.e., the system does not handle uncertainty in the UxVs' states, health, or capabilities, nor in the execution of tasks.

B. Mission Tasks

The first task type in OPS-USERS is to find targets, called the *search task*, which is a collaborative effort between the human and the automation. In OPS-USERS, the UVs are equipped with a variety of sensors for observing and searching the environment, and a target may be detected if its position lies within an UxV's sensor footprint. The automation and the human operator work in a complementary manner to accomplish efficient searching. This architecture gives the system flexibility such that the autonomy makes search decisions even in the absence of operator input, but also leverages their expertise to improve the search process, as discussed in Section II-F.

The automation generates paths for the UxVs that maximize the likelihood of detecting new targets. However, the computational complexity associated with this optimization process grows exponentially with the length of the planned paths that limits the automation to calculating localized trajectories. The human can guide the automation and prevent myopic behaviors by creating search tasks in remote, unsearched areas that are likely to contain new targets. The automated planner then assigns UxVs to the search areas the operator has identified and optimizes the local trajectories of these UxVs to efficiently search the areas of interest.

The second task type is tracking a found target, called the *track task*, which is a task supervised by the human operator, but executed by the automation. Tracking a target for a period of time decreases the uncertainty

associated with that target's state (i.e., position and velocity). However, targets are typically not tracked continuously in order to allow UxVs the flexibility of performing more tasks. When a track task ends, the tracking UxV instantiates a new track task so that the target will be revisited in the future. The revisit time is calculated based on the expected growth rate of the target's position uncertainty and the UxVs' sensing capabilities.

If a target being tracked is known to be hostile, a *neutralize hostile task* is also created. This third task type involves engaging an enemy target, and can only be performed by a weaponized UAV (WUAV). The operator must verify the classification of the enemy target as hostile and approve the weapon release.

The final task is the *refuel task*, which requires that each UxV return to base to refuel before its fuel supply is depleted. Refuel times are also scheduled by the task planner to maximize the overall mission performance, and are treated as hard constraints. UxVs may decide to refuel early if that enables them to accomplish more tasks on time. This health management task is handled entirely by the automation to reduce the operator's workload.

In order to allow the human and the automation to collaborate for task execution, the basic system architecture is divided into two major components, as shown in Fig. 2. The first is the distributed tactical planner, which is a network of onboard planning modules (OPMs) [11] that provides coordinated autonomy between the UxVs. Each UxV carries a processor that runs an instance of the OPM. The second is the ground control station, which consists of a centralized strategic planner called the central mission manager (CMM), and the operator interface (OI). These two components are discussed in turn.

C. The Distributed Tactical Planner

A decentralized implementation was chosen for the tactical planner to allow rapid reaction to changes in the environment. When appropriate, the decentralized task planner may modify the task assignment without affecting the overall plan quality (e.g., UxVs swap tasks), and it is able to make these local repairs faster through inter-UxV communication than if it had to wait for the decision to come from the human or a centralized planner. Furthermore, plans can be carried out even if the communication link with the ground control station is intermittent or lost. The architecture is scalable, since additional UxVs also add computational capability, and the decentralized framework is robust to a single point of failure, since no single UxV is globally planning for the fleet.

The decentralized task planner used in OPS-USERS is the consensus-based bundle algorithm (CBBA), a decentralized, polynomial-time, market-based protocol [12]. CBBA consists of two phases that alternate until the assignment converges. In the first phase, *task selection*, UxVs select the set of tasks for which they receive the highest reward. The UxVs place bids on the tasks they choose

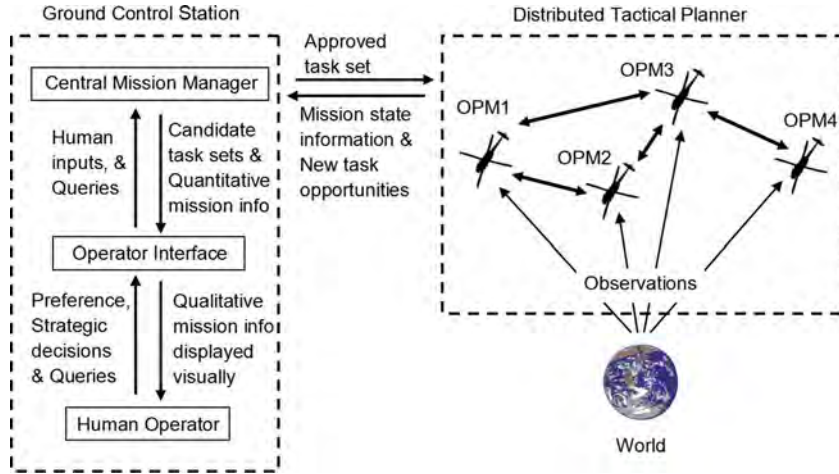


Fig. 2. OPS-USERS system architecture.

where the bids represent the marginal improvement in the score of their plan. Thus, task planning only requires the exchange of bids, which each UxV makes based on its own state only (i.e., the UxVs only need to know how much the other UxVs value the tasks they compete for—the bids—and not the locations or trajectories of the other UxVs). In the second phase, *conflict resolution*, plan information is exchanged between neighbors and tasks go to the highest bidder. CBBA is guaranteed to reach a conflict-free assignment, given a strongly connected network.

One key advantage of CBBA is its ability to solve the multiple assignment problem where each UxV is assigned a set of tasks (a *plan*), as opposed to solving the single assignment problem, where each UxV is only assigned to their next task. Planning several tasks into the future improves effectiveness in complex missions. An additional advantage is that CBBA runs in polynomial time with respect to the number of UxVs and the number of tasks. Under the implementation used for this experiment, plans could be generated for four UxVs and ten tasks in less than 1 s.

A successful distributed tactical planner must consider both task assignment as well as path planning. In OPS-USERS, the OPM generates paths in two ways. If the next task in an UxV's plan is scheduled such that the UxV must travel directly to the task location, the UxV is said to have no *spare time*. In the case where an UxV has no spare time, it calculates the minimum-time path to the task location using Dijkstra's algorithm [13]. The UxVs avoid both static obstacles and dynamic obstacles (i.e., other UxVs) along the way. It is assumed that UxVs are aware of static obstacles *a priori*, and they have an estimate of dynamic obstacles based on communicated trajectories between UxVs.

In the case where the next task in an UxV's plan is scheduled such that the UxV has more than enough time to travel to the task location, the UxV is said to have spare

time. In this case, the UxV implements a *spare time strategy*, which is chosen by the operator for each UxV before the mission begins. If the spare time strategy is to loiter, the UxV generates a trajectory to stay near their last task location until they run out of spare time and need to travel to their next task location. If the UxV's spare time strategy is to search, it generates a search trajectory that maximizes the expected probability of detecting new targets along the path. The UxV builds a breadth-first, depth-limited search tree over a receding horizon and determines which cells are observable to its sensor along each path. The UxV selects the path that has the highest cumulative probability of revealing new targets. This trajectory generation algorithm scales exponentially with the depth of the search (i.e., number of waypoints in the trajectory) and the size of the environment (i.e., number of cells in the world). For this reason, the planning horizon must be chosen such that the trajectory can be calculated at a rate at least as fast as the planning loop. For this experiment, a planning horizon of four waypoints was sufficient for the planning rate of 1 Hz. The search effort is also coordinated across the fleet to avoid redundant effort, which is known to improve search performance [14]–[16]. UxVs communicate their trajectories to each other, and UxVs receive no additional reward for observing cells, which will already be observable along another UxV's planned path.

D. Maintaining Situational Awareness

Each autonomous UxV maintains situational awareness by keeping an estimate of several mission parameters. Targets that are yet to be found are called *search targets*, and there are an unknown number of such targets at a given time. The uncertainty associated with the location of search targets is represented by a *probability map* [17]–[19], which is a discretized map of the environment. Each cell i in the map has some probability $p_i(t)$ of containing at

least one search target at time t . At each time step, the probabilities are propagated according to a target motion model. The target model quantifies the expected probability that a target moves from one cell to another cell in one time step. This transition probability is given by

$$p_{ij}^{\text{trans}} = \frac{v_{\text{tgt}} dt}{n_{\text{free}}(i)e_{ij}}.$$

The target model assumes that each target is non-evasive in that its motion is independent of the actions taken by the UxVs. It is also assumed that the probability that a target traveling from a cell i to a cell j is proportional to v_{tgt} , the target's expected speed, and dt , the time between updates. It is assumed inversely proportional to $n_{\text{free}}(i)$, the number of nonobstructed cells adjacent to cell i , and to e_{ij} , the length of the edge of cell i , in the direction of cell j . v_{tgt} is estimated in advance of a mission by a mission planner, based on the expected target types, which can include heterogeneous speeds. When a cell i is observed by an UxV at time t , its probability of containing a new target is updated according to

$$p_i(t+1) = p_i(t) \cdot (1 - f_i(t))$$

where $f_i(t)$ is the fraction of cell i that was observed at time t .

When a new target is found, it is designated a *track target*. Position, velocity, and position uncertainty estimates are kept for each track target in the environment. If a track target is not observed by any UxV at time t , the target's position uncertainty increases by a factor that is proportional to the target's last known velocity. The growth rate of the position uncertainty of a given track target is used by the planner to schedule the next desired revisit time for that target. If the uncertainty for a given target becomes too large, the target is considered lost, and becomes a search target.

The distributed tactical planner runs at a rate of one cycle per second. Each time through the planning loop, UxVs share information regarding their current position, their current planned trajectory, and updates on target position estimates. UxVs are able to update their own situational awareness based on messages from other UxVs.

E. The Ground Control Station

The ground control station is remotely connected to the UxV network and primarily consists of the CMM and the OI. The CMM is a centralized computational aid, designed to assist the operator in strategic decision making. The CMM runs a local centralized version of the task planner to determine the feasibility of assigning a given subset of tasks. The operator is able to modify the set of

tasks under consideration, and investigate a variety of "what if" scenarios, i.e., the operator changes the assigned tasks and then evaluates the impact of the new schedule before actually sending the new plan to the UxV team. The operator must approve every automation-generated plan before the tasks are released to the distributed tactical planner.

This strategic level planning does not concern the operator with the mechanics of how the tasks will be carried out (i.e., which UxV will perform which task), but allows the operator to decide which tasks should be included in the plan. This goal-based approach, as compared to individual vehicle control, is critical for single-operator control of multiple UVs since it substantially reduces operator workload, and does not require them to use valuable cognitive skills that are best reserved for knowledge-based decisions [7].

Quantitative information (i.e., UxV positions and fuel levels, target position estimates) and candidate plans are passed from the CMM to the OI, which converts this information into a visual form that is meaningful to the operator (discussed in depth in the next section). The operator can query the planner to determine the feasibility of inserting a new task into the current task list. Once the operator has decided on a plan, the approved task set is sent to the distributed tactical planner via the CMM for execution.

F. The Operator Interface

In order to allow a single operator the ability to control multiple heterogeneous UVs given the automated planner previously described, two interfaces were designed: the map display and the schedule comparison tool, detailed next.

The map display (Fig. 3) represents the primary OI, which shows both geospatial and temporal mission information (i.e., a timeline of mission events). Icons represent UxVs; low, medium, and high priority targets; and search tasks. The symbology is consistent with MIL-STD 2525 [20]. This interface also supports an instant messaging communication tool, which is a commonly used interface in military operations. This "chat" tool simulates an instant messaging tool connected to a command center that provides high-level direction and intelligence about targets in a designated area. In this experiment, this chat tool was used to communicate permission to fire orders and target priority changes, as well as issued queries to operators for mission status updates.

In the interface depicted in Fig. 3, operators' primary tasks are to identify targets found by the network of UxVs and approve weapon launches. Once a target is found, the operator is alerted to perform a target designation task (i.e., hostile, unknown, or friendly), along with assigning an associated priority level (i.e., high, medium, low). In addition, while the UxVs are capable of determining and negotiating their own search patterns, if operators are

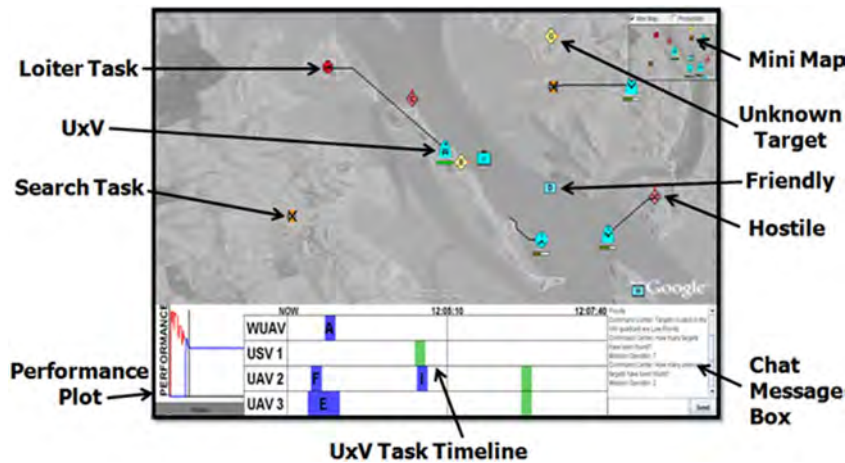


Fig. 3. Map display.

unhappy with the UxV-determined search patterns, they can create new search tasks. This occurs by operators right-clicking on those areas they feel need more attention, which brings up a menu that allows the operators to prioritize the tasks and give windows of desired coverage. This insertion of a search task by the operator in effect forces the decentralized algorithms to reallocate the UxVs such that the operators' desired tasks will be added to the task list. This human-automation interaction scheme is one of high level *goal-based* control, as opposed to more low-level *vehicle-based* control. Unlike present-day operations where multiple operators command a single vehicle, for one operator to effectively control multiple vehicles, the operator must control the network via high-level tasks as opposed to individual vehicle commands.

A performance plot in the lower left corner of Fig. 3 gives operators insight into the automated planner performance as the graph shows expected performance (given an *a priori* cost function) compared to actual performance. When the automation generates a new plan with a predicted performance that is 10% higher than the current plan's performance, the *replan* button in Fig. 3 turns green and flashes. This illuminated *replan* button indicates that a new plan is ready for operator approval, and when the *replan* button is selected, the operator is taken to the schedule comparison tool (SCT, Fig. 4).

The three geometrical forms at the top of the SCT in Fig. 4 are configurational displays that enable the operator to quickly compare three schedules: the current, working, and proposed schedules. Configurational displays allow

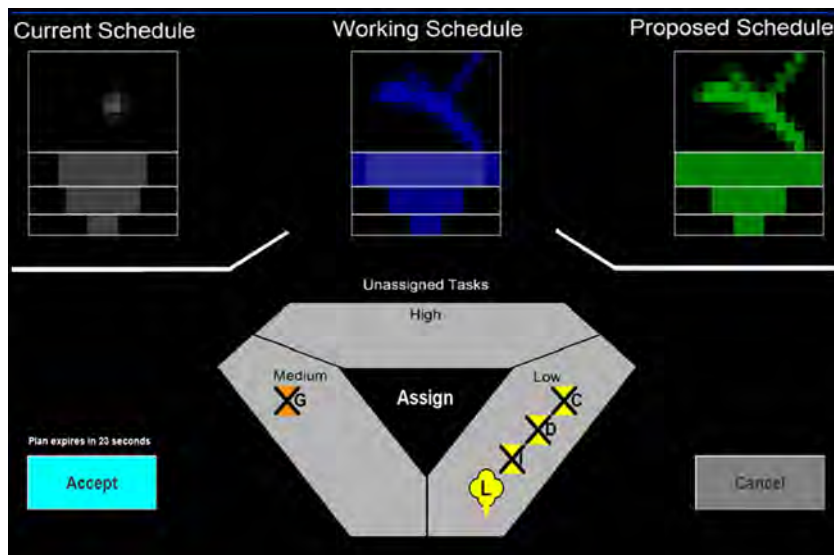


Fig. 4. The schedule comparison tool.

operators the ability to utilize more efficient perceptual processes rather than cognitively demanding processes that rely on memory and inference [21]. This reliance on perceptual reasoning is known as direct perception, and for this decision support display in Fig. 4, the goal is to graphically depict to the operator how much area will be searched given the plan, as well as how many tasks of the differing priorities will be accomplished.

In Fig. 4, the left form (gray) is the current UxV schedule. The right form (green) is the latest automation proposed schedule. The middle schedule (blue) is the working schedule that results from the user modifying the plan by querying the automation to assign particular tasks. The rectangular grid on the upper half of each shape represents the estimated area that the UxVs will search according to the plan. The hierarchical priority ladders show the percentage of tasks assigned in high (top), medium (middle), and low (bottom) priority levels. Thus, in keeping with the direct-perception paradigm, the shape that is filled with the most color depicts the best plan. In addition, while the SCT button turns green whenever the automation determines a 10% better plan, operators can select this button at anytime to rearrange the task list if desired.

When the operator first enters the SCT, the working schedule is identical to the proposed schedule. The operator can conduct a “what if” query process by dragging the desired unassigned tasks into the large center triangle in Fig. 4 labeled “assign.” This query forces the automation to generate a new plan if possible, which then becomes the working schedule. The configural display of the working schedule alters to reflect the changes of newly assigned tasks. However, due to possible resource shortages, all tasks may not be assigned to the UxVs, which is representative of real-world constraints. When a task cannot be assigned, the task that was dragged into the center triangle pops out, representing the case where the automated planners could not honor the operator’s requests due to a constraint violation in time, fuel required, or priority conflict (i.e., only tasks with higher priority could be accommodated).

The working schedule configural display updates with every individual query so that the operator can quickly compare the three schedules based on quantity of color. This querying process represents a more collaborative effort between the human and automation, which has been shown to improve operator performance and situation awareness in similar complex settings [22]. Ultimately, the operator either accepts the working or proposed schedule or can cancel to continue with the current schedule.

The system is designed so that either both screens can be displayed if the operator’s ground control stations support such screen real estate, or the screens alternate if only a single display is available. The single display paradigm was used in the experiment detailed in the next section since the expected environment was one of an expeditionary controller, i.e., a human moving through the environment with only a laptop to control the UxVs.

III. THE EXPERIMENT

In order to determine the impact of human interaction with the decentralized vehicle network on mission performance, a human-on-the-loop experiment was conducted and the results compared to those of a perfectly obedient human who always agreed with the automated planner’s proposed schedule. For this experiment, participants were responsible for one rotary-wing UAV, one fixed wing UAV, one unmanned surface vehicle (USV), and a WUAV. The UAVs and USVs were responsible for searching for targets, which each had a unique ID internal to the system so ground truth was always known to the experimenters for later analysis. The targets continually moved on predetermined paths (unknown to the UxVs and operators), so there was no guarantee that the targets would be found. Once designated, hostile targets were tracked by one or more UxVs until the human operator approved WUAV missile launches. For the purposes of this experiment, it was assumed that the UxVs could not be destroyed.

The experiment was conducted on a Dell Optiplex GX280 with a Pentium 4 processor and an Appian Jeronimo Pro 4-Port graphics card. The display’s resolution was 1280×1024 pixels. To familiarize each subject with the interfaces in Figs. 1 and 2, a self-paced, slide-based tutorial was provided, as well as a practice session. The training took approximately 30 min to complete. Practice was followed by three 10-min test sessions, representing each of the three possible replan intervals, discussed in detail in the next session, in a counterbalanced and randomized order. Each scenario was different, but similar in difficulty. The interface recorded all operator actions.

Because it has been previously established that the rate at which an automated scheduling tool prompts the user for intervention can negatively influence UAV operator performance due to high workload [22], the subjects in this experiment were all exposed to three different replanning intervals of 30, 45, and 120 s. This means that the operator was prompted by the green illumination of the replan button and an aural replan alert every 30/45/120 s, indicating an improved schedule was available. In actual use, such automated planners would not likely generate plans with such regularity. However, the rate of operator prompting was fixed at these intervals (with a standard deviation of ± 5 s) for experimental control. OPS-USERS can typically generate a new, theoretically improved plan once per every 10 s, so these new plans were simply suppressed until the correct window for that session replan interval was achieved.

The original subject population consisted of 31 subjects. However, analysis of the resulting data through statistical clustering revealed that only approximately one-third of the subjects ($N = 9$) actually followed instructions to collaborate with the automation and evaluate the automation’s plans at the 30/45/120 s intervals. These subjects

were labeled consenters, in that they consented to the replanning schedule, although they did not necessarily agree to the proposed plan, i.e., the consenters changed the automation's proposed schedule about one-third of the time. Consenters had statistically significantly improved performance over the remaining 22 subjects who were labeled dissenters (who essentially ignored the automation) or mixed consenters (who only sometimes agreed with the automation) in terms of when to replan.

Characteristics of the consenters were higher than average video gaming, as well as no military experience, i.e., those participants with military experience had negative attitudes towards UVs. An in-depth analysis of the consenter versus dissenter performance is detailed elsewhere [23], but for the purposes of this paper which assesses the added value of human-automation collaboration, only the data from the consenter group will be considered since they were statistically superior performers as compared to the remaining 22 participants. For the remainder of this paper, this consenting group will be called collaborators since they generally agreed to work with the automation, with occasional input to redirect the automation as necessary. They did not necessarily agree with the automation, on average they disagreed with 30% of automation-generated plans and conducted "what if" queries to generate their own slightly modified plans.

After the human-on-the-loop experiment, the second portion of this experiment was conducted during which a single operator played the role of a perfectly obedient operator. This operator completed a total of 27 counter-balanced trials (nine each at the 30/45/120 s intervals). This operator always agreed with the replanning rates, i.e., always entered the SCT when prompted, and also agreed with the proposed schedule, i.e., the operator never attempted to change the plan. In addition, this operator never inserted new unprompted search tasks, in effect always trusting that the automation was performing the optimal search and tracking task allocation. The use of a single operator reduced variability in response time, which was generally less than 500 ms, typical for such point-and-click behavior that requires no decision making on the part of the operator.

Performance-dependent variables for both phases of the experiment included the percentage of area covered (which is critical since one primary task of the UxVs was to cover as much area as possible), number of targets found, number of hostile targets neutralized, and the average operator utilization for a given session. Utilization, a workload measure, is the percent time an operator is busy over the entire mission. Operators were considered busy when performing one or more of the following tasks: creating search tasks, identifying and designating targets, approving weapons launches, interacting via the chat box, and replanning in the SCT. Monitoring time was not included in percent busy time. Previous research has shown that in single-operator goal-based control of multiple

entities, operator performance can significantly drop when tasked greater than 70% in terms of utilization [24], [25].

IV. RESULTS AND DISCUSSION

Since this mission was one of searching, tracking found targets, and neutralizing hostile targets, each of these mission elements constitutes an effects-based performance metric. These are now discussed in turn. For all statistical tests reported, $\alpha = 0.05$.

The first mission performance metric, percentage of area searched (Fig. 5), illustrates the effectiveness of the overall system when operators collaborated with the algorithm (which meant changing the algorithm's plan approximately 30% of the time), as compared to when the human was perfectly obedient and always executed the automation's plan. The overall mixed 3×2 ANOVA (with repeated measures across the three replanning intervals) shows a statistically significant difference between collaborative and obedient humans [$F(1, 24) = 13.3, p < 0.001$]. There was no overall main effect for the replanning interval [$F(2, 24) = 1.7, p = 0.171$] or interaction between replanning interval and obedient versus collaborative human [$F(2, 24) = 0.711, p = 0.501$]. Bonferroni pairwise comparisons reveal that for the 30- and 45-s conditions, the collaborative human-automation team is superior to the obedient human ($p < 0.001$ for both comparisons), but that there was no statistical difference at the 120-s interval ($p = 0.411$). This indicates that for the faster replan intervals, the collaborative human provided significant advantage to the algorithm, increasing the overall area covered by up to 30%, on average. However, this large gap closed in the longer replan intervals, with the collaborative




Fig. 5. Search task performance.




Fig. 6. Target-finding task performance.

human not appreciably improving the percent area covered compared to what the automation essentially planned on its own.

It should also be noted that the collaborative human effort for searching in terms of area covered was statistically no different between the 45- and 120-s replanning intervals ($p = 0.355$), but significantly lower at the 30-s interval ($p = 0.002$, as compared to 120 s). However, the automation's performance with respect to search linearly increased as replanning interval increased. The reason for this improvement is likely due to the fact that the UxVs spend a greater amount of time searching when the replanning interval is longer. Search is a spare time strategy, and so if the replan interval is longer, other tasks such as track or neutralize may be delayed because they are not released to the distributed tactical planner until the next replan. These results suggest that there was a plateau for human effort, but due to the limited experimental conditions, it is not evident in the data that the automation reached a plateau. This will be discussed further in a subsequent section.

Related to the area searched metric is the targets found metric, since simply searching the most space is not a comprehensive performance metric, i.e., it represents a subset of objectives. Since ultimately the goal was to destroy as many targets as possible, the number of targets found can be a proxy measure for the quality of search. Using the same 3×2 ANOVA model discussed previously, the metric of percentage of available targets found by participants (Fig. 6) demonstrates that when the humans worked with the planner, significantly more targets were found across all replanning intervals ($F(1, 24) = 17.8, p <$

0.001), than when the human perfectly obeyed automation directives. Again, there were no other main effects for replanning rate [$F(2, 24) = 0.470, p = 0.628$] or interaction [$F(2, 24) = 1.85, p = 0.179$].

With humans occasionally redirecting the automation (i.e., changing the automation's proposed schedule 30% of the time), up to 50% improvement was seen in the number of targets found, as compared to if the automation had been left alone. Unlike the area-searched metric, there was no convergence of human-automation performance in that the human-automation team always performed more than 20% better for every replanning rate condition. Once again, the collaborative human performance appeared to plateau for the 45- and 120-s intervals in that they were statistically no different ($p = 0.891$), while the 30-s condition was again statistically different from the 120-s condition ($p = 0.046$).

Another important mission performance metric was the percentage of time targets were successfully tracked, i.e., not lost. Recall that targets were not continually tracked, in that UxVs could multitask, and attempt to track multiple targets. In this experiment, we were not able to detect any measurable human operator input that improved upon the automation, which was able to effectively track targets 79%–99% of the time. This is not to say that the distributed planner would perform this well in all cases, but just for the numbers of targets and UxVs in this short experimental scenario, the planner performed very well and did not significantly lose target tracks. Moreover, since human operators could only assist help with finding lost targets, they could only indirectly influence tracking targets. However, how often operators inserted and reprioritized search tasks could affect tracking performance, and we leave this as an area of future research.

For the last mission performance metric, hostiles neutralized, there was no statistical difference between the collaborative and obedient humans, meaning that active human management of this process did not improve upon the automation's performance [$F(1, 24) = 0.001, p = 0.974$]. In addition, there was no statistical difference in the replanning intervals [$F(2, 24) = 1.36, p = 0.272$], so that regardless of the 30-, 45-, or 120-s interval, the hostiles neutralized performance was consistent. There was no significant interaction [$F(2, 24) = 0.400, p = 0.675$]. These results are not surprising, given that of the three mission performance requirements, human intervention for hostile neutralization was minimal, i.e., once a target was designated as hostile, execution of the neutralization task was very rule based, in that as soon as a target was marked hostile, it was shadowed by a WUAV until an order was received to neutralize. Given the clearly established rule set to manage this process, the automation was very capable of positioning the WUAV correctly, without human intervention.

Given that workload is a significant concern in single-operator control of multiple UxVs, operator

utilization (i.e., percent busy time) was measured across the three replanning intervals. The results showed no significant difference in utilization across the three replanning levels [$F(2, 24) = 2.817, p = 0.080$], which is not surprising since the operators in this analysis represent the top performers. The average percent utilization was 44% ($\pm 7\%$), indicating that operators were not overloaded and well below the threshold of 70%.

A. Function Allocation Analysis

The results from the experiment clearly show that ultimately mission success was dependent on the collaboration between the operators and the automation, and that mutually exclusive function allocation either within or across tasks is not an effective solution. Given the mission for this team of heterogeneous UVs of search, track, and neutralize enemy targets and the results seen from this experiment, Table 1 summarizes the function allocation between the human and the automation for the various mission components that promoted the best overall system performance, as motivated by the statistical results in the previous section. As Figs. 5 and 6 demonstrate, the best search performance was obtained when the human guided the automation, as opposed to leaving the search task to just the automation. These results are somewhat expected, as the human ability to conduct relatively effortless effective visual searches as compared to computers has long been recognized [26]. However, these results demonstrate just how much advantage human guidance in the search task can provide. In addition, these results suggest that human performance could plateau as a function of increased replan times, but the results are conflicting in terms of the capabilities of the distributed planners.

For this set of experimental conditions, human performance in terms of the search task appeared to peak at the 45-s replanning interval, with no statistically significant improvement when replanning at the 120-s interval. For the percent area covered, the automated planner did worse than when augmented with human guidance at the 30- and 45-s intervals, but at the 120-s interval, it performed the same. This is possible evidence that there may have been a replanning interval that would have promoted better automated planning performance in the search task. This is an area of current active research, as it is not obvious how to determine, given an uncertain environment with emergent targets, how often a distributed network of

UVs should communicate in order to maximize the search effort.

While the benefit of the human ability to aid in the search task appeared to become on par with that of the automation as the time intervals between planning grew longer, this relationship did not hold true for the targets found metric. In this case, the human-automation team far exceeded the automation's ability to find targets. While percent of area searched is an important metric given the stated mission, number of targets found demonstrates the effectiveness or quality of the search. Thus, humans, aided by the automation, were better able to judge not just where the automation should be searching in terms of areas left uncovered, but also when areas should be revisited and when to change a course of action when a current plan appeared to be suboptimal. It is this combination of recognizing both where and when to search that led to superior human-automation teaming in the target-finding task.

While some tasks benefitted from a human-automation collaborative effort, there are some tasks that should be primarily executed by either just the UxVs or humans. For example, given the severe consequences of misidentification of a target and the lack of technical progress that has been made in automated target recognition, target identification remains a task that is expected to remain a human endeavor for the near future. However, improved sensor technology may be able to augment and aid the human in this complex, high-risk task.

Alternatively, tracking performance did not improve with human assistance, suggesting that this task was best left to automated planners. However, one caveat is that a large range of possible scenarios was not examined in this experiment, so under different conditions, results could have varied. However, the tracking task is one of complex, rapid computation in that several variables must be considered such as locations and likely trajectories of targets, locations, and movement tracks of UxVs, and the need to revisit targets to maximize tracking coverage in a limited resource environment. Given the inability of a single operator to be able to make such rapid, precise computations in a dynamic environment, it is expected that automation would be able to execute this task assuming that the target was clearly identified. Recent developments in automated target tracking have been realized in operational settings such that, for example, a designated target in the form of a white truck can be automatically tracked as it moves. It

Table 1 Agent(s) That Promoted the Best Task Performance in the Human-on-the-Loop Experiment

		Agent (s)		
		Automation Led	Human + Automation	Human Led
Task	Search		✓	
	Target Identification			✓
	Target Track	✓		
	Neutralize Hostiles		✓	

remains to be determined how humans could augment or assist in these real-world tracking tasks.

Last, it should be noted that the hostile neutralization task was another that benefitted from a collaborative human-automated planner effort. This task is a collaborative one because the human operator made the final decision to release a weapon, which is a stressful and information-gathering intensive task. The automation kept the WUAV in place so that if and when a command for neutralization was received, the system was ready to respond. Thus, the operator was able to offload lower level position keeping tasks to the automation, while reserving precious cognitive resources for an overall knowledge-based task of firing a weapon that will likely have grave consequences.

V. CONCLUSION

In the future concept of one operator supervising multiple collaborative UxVs, the potential exists for high operator workload and negative performance consequences. As a result, significant autonomy is needed to aid the operator in this multiple UxV control task. Due to the dynamic and uncertain nature of the environment, control of collaborative and decentralized UxVs requires rapid automated replanning. However, as demonstrated in this study, human management of the automated planners is critical, as automated planners cannot always generate accurate solutions for every combination of events. Though fast and able to handle complex computation far better than humans, computer optimization algorithms are often “brittle” in that they can only take into account those quantifiable variables identified as critical during the design stage [27], [28].

This research has shown, for an admittedly narrow set of conditions, that given a decentralized UV network, human guidance can provide substantial benefit in search tasks, with arguable little improvement in tracking tasks. Most importantly, this research has demonstrated that there is a shared space in such missions for human-automation collaboration. While there may be some mission tasks that can be mutually assigned to either a human or UxV (due primarily to inherent limitations), tasks exist that improve substantially when humans and automation work together. Similar results have been achieved in human-guided optimization of heuristic search algorithms

in other scheduling domains [29], as well as in single-operator supervisory control of a swarm of UAVs for target pursuit [30]. Current work is underway to identify characteristics of scheduling and planning problems and algorithms that make them particularly suitable for human-algorithm collaboration.

This research also raises a set of new questions that deserve further consideration. While algorithms such as those embedded in OPS-USERS can typically generate new schedules in just seconds, often these schedules are not optimal, i.e., solutions generated by the predetermined static cost function may not reflect the true state of all variables in a dynamic command and control scenario. Such planners will consider even a solution that is 1% better than a previous solution as a “better solution,” but often such algorithms cannot account for “satisficing” behavior [31] that ultimately is just as “optimal” (from the viewpoint of stakeholders). Reducing the gap between planner assumptions and human expectations is an important area of future research. Moreover, in this study, allowing more time between replans improved the search performance metrics, so future work is needed to objectively determine effective replanning intervals. These intervals are likely dynamic parameters, and more research is needed to determine if some Pareto front exists that could predict a region of optimized performance.

However, changing the replan interval to better suit planner performance could introduce new, unintended consequences. It remains an open area of research to determine how to design such a system that achieves a balance between human guidance and human interference. Previous research has shown that those operators that refuse to collaborate with the automation can actually degrade overall system performance and also significantly increase their own workload [23]. Thus, designing such a system that reduces operator workload and engenders trust and collaboration is critical. ■

Acknowledgment

The authors would like to thank A. Clare, C. Hart, and D. Southern (Massachusetts Institute of Technology (MIT), Cambridge) and P. Maere (Delft University of Technology, Delft, The Netherlands) who were instrumental in gathering the data for this effort.

REFERENCES

- [1] *World Unmanned Aerial Vehicle Systems 2010*, Teal Group Corporation, Manassas, VA, 2009.
- [2] R. Teo, J. S. Jang, and C. Tomblin, “Automated multiple UAV flight—The Stanford dragon fly UAV program,” in *Proc. 43rd IEEE Conf. Decision Control*, Nassau, Bahamas, 2004, DOI: 10.1109/CDC.2004.1429422.
- [3] M. Alighanbari and J. How, “Decentralized task assignment for unmanned aerial vehicles,” in *Proc. IEEE Conf. Decision Control/Eur. Control Conf.*, Sevilla, Spain, 2005, pp. 5668–5673.
- [4] J. Tisdale, A. Ryan, K. Zu, D. Tornqvist, and J. K. Hedrick, “A multiple UAV system for vision-based search and localization,” in *Proc. Amer. Control Conf.*, 2008, pp. 1985–1990.
- [5] N. Nigam, S. Bieniawski, I. Kroo, and J. Vian, “Control of multiple UAVs for persistent surveillance: Algorithm description and hardware demonstration,” in *Proc. AIAA Infotech@Aerospace Conf.*, Seattle, WA, 2009.
- [6] *Network Centric Warfare: Department of Defense Report to Congress*, Department of Defense, Office of the Secretary of Defense, Washington, DC, 2001.
- [7] M. L. Cummings, S. Bruni, S. Mercier, and P. J. Mitchell, “Automation architecture for single operator-multiple UAV command and control,” *Int. Command Control J.*, vol. 1, pp. 1–24, 2007.
- [8] M. L. Cummings and P. J. Mitchell, “Predicting controller capacity in supervisory control of multiple UAVs,” *IEEE Trans.*

- Syst. Man Cybern. A, Syst. Humans, vol. 38, no. 2, pp. 451–460, Mar. 2008.
- [9] S. Dixon, C. Wickens, and D. Chang, “Mission control of multiple unmanned aerial vehicles: A workload analysis,” *Human Factors*, vol. 47, pp. 479–487, 2005.
 - [10] M. A. Goodrich, “On maximizing fan-out: Towards controlling multiple unmanned vehicles,” in *Human-Robot Interactions in Future Military Operations*, M. Barnes and F. Jentsch, Eds. Surrey, U.K.: Ashgate, 2010.
 - [11] J. P. How, C. Fraser, K. C. Kulling, L. F. Bertuccelli, O. Toupet, L. Brunet, A. Bachrach, and N. Roy, “Increasing autonomy of UAVs decentralized CSAT mission management algorithm,” *IEEE Robot. Autom.*, vol. 16, no. 2, pp. 43–51, Jun. 2009.
 - [12] H. L. Choi, L. Brunet, and J. P. How, “Consensus-based decentralized auctions for robust task allocation,” *IEEE Trans. Robot.*, vol. 25, no. 4, pp. 912–926, Aug. 2009.
 - [13] E. W. Dijkstra, “A note on two problems in connexion with graphs,” *Numerische Mathematik*, vol. 1, pp. 269–271, 1959.
 - [14] M. Flint, M. Polycarpou, and E. Fernandez-Gaucherand, “Cooperative path-planning for autonomous vehicles using dynamic programming,” in *Proc. 15th Triennial World Congr.*, Barcelona, Spain, 2002, DOI: 10.3182/20020721-6-ES-1901.01305.
 - [15] I. Maza and A. Ollero, “Multiple UAV cooperative searching operation using polygon area decomposition and efficient coverage algorithms,” *Distrib. Autonom. Robot. Syst.*, vol. 6, pp. 221–230, 2007.
 - [16] K. R. Guruprasad and D. Ghose, “Multi-agent search using Voronoi partitions,” in *Proc. Int. Conf. Adv. Control Optim. Dyn. Syst.*, Bangalore, India, 2007, pp. 380–383.
 - [17] F. Bourgault, T. Furukawa, and H. F. Durrant-Whyte, “Optimal search for a lost target in a Bayesian world,” *Field Service Robot.*, vol. 24, pp. 209–222, 2003.
 - [18] Y. Jin, A. Minai, and M. Polycarpou, “Cooperative real-time search and task allocation in UAV teams,” in *Proc. IEEE Conf. Decision Control*, 2003, vol. 1, pp. 7–12.
 - [19] M. Flint, E. Fernandez-Gaucherand, and M. Polycarpou, “A probabilistic frame-work for passive cooperation among UAVs performing a search,” in *Proc. 16th Int. Symp. Math. Theory Netw. Syst.*, 2004.
 - [20] *Interface Standard, Common Warfighting Symbolology MIL-STD-252B*, Department of Defense, Washington, DC, Jan. 1999.
 - [21] J. J. Gibson, *The Ecological Approach to Visual Perception*. Boston, MA: Houghton Mifflin, 1979.
 - [22] M. L. Cummings, A. S. Brzezinski, and J. D. Lee, “Operator performance and intelligent aiding in unmanned aerial vehicle scheduling,” *IEEE Intell. Syst.*, vol. 22, *Special Issue on Interacting With Autonomy*, no. 2, pp. 52–59, Mar./Apr. 2007.
 - [23] M. L. Cummings, A. Clare, and C. Hart, “The role of human-automation consensus in multiple unmanned vehicle scheduling,” *Human Factors*, vol. 52, pp. 17–27, 2010.
 - [24] D. K. Schmidt, “A Queuing analysis of the air traffic controller’s workload,” *IEEE Trans. Syst. Man Cybern.*, vol. SMC-8, no. 6, pp. 492–498, Jun. 1978.
 - [25] M. L. Cummings and C. Nehme, “Modeling the impact of workload in network centric supervisory control settings,” in *Neurocognitive and Physiological Factors During High-Tempo Operations*, S. Kornuth, Ed. Surrey, U.K.: Ashgate, 2010, pp. 23–40.
 - [26] P. M. Fitts, *Human Engineering for an Effective Air Navigation and Traffic Control System*. Washington, DC: Nat. Res. Council, 1951.
 - [27] P. Smith, E. McCoy, and C. Layton, “Brittleness in the design of cooperative problem-solving systems: The effects on user performance,” *IEEE Trans. Syst. Man Cybern. A, Syst. Humans*, vol. 27, no. 3, pp. 360–371, May 1997.
 - [28] B. G. Silverman, “Human-computer collaboration,” *Human-Computer Interaction*, vol. 7, pp. 165–196, 1992.
 - [29] G. W. Klau, N. Lesh, J. Marks, and M. Mitzenmacher, “Human-guided search,” *J. Heuristics*, vol. 16, no. 3, pp. 289–310, 2010.
 - [30] F. Legras and G. Coppin, “Autonomy spectrum and performance perception issues in swarm supervisory control,” *Proc. IEEE*, 2012, DOI: 10.1109/PROC.2011.2174103.
 - [31] H. A. Simon, R. Hogarth, C. R. Piott, H. Raiffa, K. A. Schelling, R. Thaler, A. Tversky, and S. Winter, *Decision Making and Problem Solving*. Washington, DC: Nat. Acad. Press, 1986.

ABOUT THE AUTHORS

M. L. Cummings (Senior Member, IEEE) received the Ph.D. degree in systems engineering from the University of Virginia, Charlottesville, in 2004.

She is an Associate Professor in the Aeronautics and Astronautics Department and the Director of the Humans and Automation Laboratory at the Massachusetts Institute of Technology (MIT), Cambridge.

Dr. Cummings is an Associate Fellow of the American Institute of Aeronautics and Astronautics (AIAA), a member of the National Research Council Board on Human Systems Integration, and a member of the Naval Research Advisory Committee.



Jonathan P. How (Senior Member, IEEE) received the B.A.Sc. degree from the University of Toronto, Toronto, ON, Canada, in 1987 and the S.M. and Ph.D. degrees in aeronautics and astronautics from the Massachusetts Institute of Technology (MIT), Cambridge, in 1990 and 1993, respectively.

He then studied for two years at MIT as a Postdoctoral Associate for the Middeck Active Control Experiment (MACE) that flew onboard the Space Shuttle Endeavour in March 1995. Currently, he is the Richard Cockburn MacLaurin Professor of Aeronautics and Astronautics at MIT. Prior to joining MIT in 2000, he was an Assistant Professor in the Department of Aeronautics and Astronautics, Stanford University, Stanford, CA.



Prof. How was the planning and control lead for the MIT DARPA Urban Challenge team that placed fourth. He was the recipient of the 2002 Institute of Navigation Burka Award and a recipient of a Boeing Special Invention award in 2008. He is an Associate Fellow of the American Institute of Aeronautics and Astronautics (AIAA).

Andrew Whitten received the B.S.E. degree in aerospace engineering from the University of Michigan, East Lansing, in 2008 and the S.M. degree from the Department of Aeronautics and Astronautics, Massachusetts Institute of Technology (MIT), Cambridge, in 2010.

He is a Performance and Flying Qualities Engineer for the United States Air Force. He works for the Global Vigilance Combined Test Force, Edwards Air Force Base (AFB), CA, flight testing the RQ-4 Global Hawk UAV.



Olivier Toupet received the M.S. degree in aerospace engineering from SUPAERO, Toulouse, France, in 2004 and the M.S. degree in aeronautics and astronautics from the Massachusetts Institute of Technology (MIT), Cambridge, in 2006.

He is an expert in multivehicle command and control and optimization and, at the time of this research effort, was a lead Autonomy, Controls and Estimation (ACE) Research Engineer at Aurora’s Research and Development Center, Cambridge, MA.

